

UNITED STATES PATENT APPLICATION

of

THOMAS B. BOLT

for

EMULATED BACKUP TAPE DRIVE USING DATA COMPRESSION

FIELD OF THE INVENTION

10 The present invention relates generally to back up data storage, and more specifically, to an emulated backup tape drive that stores non-compressed data during backup operations and then afterwards when the drive is idle, retrieves, compresses and then re-stores the data to reclaim space on the storage medium of the drive.

BACKGROUND

15 With the increasing popularity of Internet commerce and network centric computing, businesses and other entities are becoming more and more reliant on information. Protecting critical data from loss due to system crashes, virus attacks and the like is therefore of primary importance. A well designed data protection program will generally have the ability to (i) instantly re-store data in the event of a disaster to enabled continued computing operations; (ii) re-store
20 data over an extended period of time (hours or days) without disrupting normal computing operations; and (iii) archive copies of data that are retrieved infrequently and with little urgency. Tape drives have long been a choice for storing archival back up data in information systems.

25 Historically many such tape drives have used data compression to maximize the amount of data that can be stored on the tape. Tape, however, is

2050220 020502 007243 007243

a relatively slow and inefficient storage medium. Consequently emulated "tape" drives that use arrays of hard drives have become more popular recently. These emulated tape drives often rely on data compression to enable the storage of more data. The problem with current emulated tape drives is that the data compression is performed "on the fly" during the backup. In other words, compression occurs in the critical path of the down loading of data, thereby impeding performance. The designers of emulated tape drive systems have therefore relied on expensive, high speed, hardware data compression solutions to achieve an acceptable level of performance. The use of slower, less expensive software compression algorithms have not been a viable option in the past because of a lack of acceptable performance.

An emulated backup tape drive that stores non-compressed data during backup operations and then afterwards when the drive is idle, retrieves, compresses and then re-stores the data to reclaim space on the storage medium of the drive is therefore needed.

SUMMARY

To achieve the foregoing, and in accordance with the purpose of the present invention, a back up storage device is disclosed that stores non-compressed data during backup operations and then afterwards when the device is idle, retrieves, compresses and re-stores the data to reclaim space on the storage medium of the device. During operation, a duty cycle having a backup window period and an idle period is defined. When back ups occur during the window, data is down-loaded and stored on the device in non-compressed form. Later during the idle period of the duty cycle, the non-compressed data is retrieved, compressed and re-stored to reclaim space on the storage medium of the device. Since the compression occurs when the back up device is idle, the rate at which data is backed up is not adversely effected in any way. Thus a low cost software data compression algorithm may be used. In one embodiment of the invention, the back up storage device is an emulated tape drive that uses an array of hard drives for the storage medium.

In other embodiments, any type of storage medium can be used such as tape or semiconductor memory chips for example.

BRIEF DESCRIPTION OF THE DRAWINGS

5 The invention, together with further advantages thereof, may best be understood by reference to the following description taken in conjunction with the accompanying drawings in which:

Figure 1 is block diagram of an exemplary information infrastructure in which the emulated backup tape drive of the present invention may be used.

10 Figures 2A, 2B and 2C are diagrams of the emulated tape drive of the present invention.

Figure 3 is block diagram of a controller of the emulated tape drive of the present invention.

Figure 4 is flow diagram illustrating a backup and compression duty cycle of the emulated tape drive of the present invention.

DESCRIPTION

15 Referring to Figure 1, a block diagram of an exemplary information infrastructure in which the emulated backup tape drive of the present invention may be used is shown. The information infrastructure 10 includes a plurality of
20 clients 12 and a plurality of servers 14 coupled together by a client network 16, a primary storage location 18, and one or more emulated tape drives 20 coupled together by a network connection 22. The clients 12 can be any type of client such as but not limited to a personal computer, a "thin" client, a personal digital assistant, a web enabled appliance or a cell phone. The servers 14 may
25 also be of any variety such as those based on the Unix, Linux, or the Microsoft Windows operating systems or a combination thereof. Likewise, the client network 16 can be any type of network including but not limited to the Internet, a corporate intranet, a wide area network, a local area network, a wireless network, or any combination thereof. The primary storage location 18 can be

arranged in a number of different types of configurations, such as a storage array network (SAN), or network attached storage (NAS), or direct attached storage . In other embodiments, the primary storage location 18 can reside in the chassis/cabinet of the servers 14, stand alone storage devices, or a combination thereof. The connection 22 can be either direct (such as parallel SCSI or IDE) or a network topology such as fibre channel, Ethernet (fast, gigabit, or 10 gigabit) for example. Also when multiple emulated tape devices 20 are used, they may be daisy-chained together to provide more backup capacity.

10 Referring to Figure 2A, a perspective view of an emulated tape drive 20 is shown. The emulated tape drive 20 includes a chassis 30, a power supply 32, a pair of fans 34, and an array of hard drives 36, all housed in the chassis 30. A user interface panel 38 is located at the front of the chassis 30. A back plate 40 is provided at the rear portion of the chassis 30.

15 Referring to Figure 2B, an exploded perspective view of the hard drives 36 according to one embodiment of the present invention is shown. In this embodiment, the hard drives 36 are configured in a pair of left and right rails (42L and 42R) respectively. Within each rail 42, five (5) columns of three (3) hard drives 36 are arranged in disk packs 44 respectively. It should be noted
20 that this embodiment is only exemplary. According to various other embodiments of the present invention, the number of hard drives 36 may be arranged in any number of rails 42 (rows) and the number of hard drives 36 (columns) per disk pack 44 may vary.

Referring to Figure 2C, a view of the back plate 40 of the chassis 30 is
25 shown. The back plate 40 includes vents for the fans 34 and a number of input/output ports 46. The input/output ports 46 are provided to connect a controller 48 (not visible because it is internal to the chassis 30) to the primary storage location 18 through the network connection 22. For more details on the features and operation of the emulated tape drive 20, see co-pending US
30 application entitled "Storage System Utilizing An Active Subset of Drives During Data Storage And Retrieval Operations" by Thomas B. Bolt and Kevin C. Daly,

attorney docket no. Q02-1037.US1, filed concurrently herewith incorporated herein by reference for all purposes.

Referring to Figure 3, a block diagram of one embodiment of the controller 48 configured for the the emulated tape drive 20 illustrated in Figures 2A-2C is shown. With this embodiment, the controller 48 includes a micro-controller 50, such as a microprocessor, configured to communicate with the hard drives 36 of the disk packs 44 through a USB controller 52, a USB hub 54 and bridge circuit 56. For the sake of simplicity, these components are shown for only one disk pack 44. The remaining four disk packs 44 of the right rail 42R and all of the disk packs 44 of the left rail 42L communicate with the micro-controller 50 in a similar arrangement. In situations where the network connection 22 is fiber channel, the micro-controller 50 is connected to the primary storage location 18 through an optical transceiver 58 and a fiber channel controller 60. Alternatively, when the network connection 22 is a Gigabit Fast ethernet connection, the micro-controller 50 is connected to the primary storage location 18 through an ethernet transceiver 62 and an ethernet controller 64. It should be pointed out that these two connections are merely illustrative. In various embodiments of the present invention, multiple fiber channel ports and/or multiple ethernet channel ports, either alone or in any combination, can be provided. Alternate interfaces such as parallel SCSI (Small Computer System Interface) may also be substituted or used in conjunction with fire channel of Ethernet. Additionally, alternate internal interconnect technologies such as fibre channel or parallel SCSI could be used instead of USB.

A system memory 66 and a non-volatile memory 68 are also coupled to the micro-controller 50. In one embodiment of the invention, the system memory 66 is RAM and the non-volatile memory is Flash. The non-volatile memory is used for storing the micro-code used to program the micro-controller 50 as well as the compression/decompression software algorithms which are used by the emulated tape drive 22. It again should be noted that the circuit components of this diagram are merely illustrative of one embodiment of the present invention. Other embodiments would be readily apparent to those

skilled in the art. For example, in embodiments with either more or fewer rails 42 and disk packs 44, additional or fewer USB controllers 52 and USB hubs 54 would be required. Also the interface hardware between the input/output ports 44 would be different if other type of networking protocols besides fiber channel or ethernet were used.

Referring to Figure 4, a flow diagram 80 illustrating the operation of the emulated tape drive 20 is shown. Initially a system administrator or user of the information infrastructure 10 defines a duty cycle (step 82) for the emulated tape drive 20. The duty cycle includes a backup window period and an idle period. Usually the backup window period is scheduled at a set time each day or at some other fixed time interval when the emulated tape drive 20 is at its lowest utilization. When backups are not occurring, the emulated tape drive 20 is idle.

According to one embodiment, when the duty cycle starts (step 84), the backup of data begins (step 86). The data is downloaded from the primary storage location 18 through the input/output ports 44, the micro-controller 50, the appropriate USB controller 52, USB hub 54, and bridge circuit 56 and stored (step 88) in non-compressed form on one of the hard drives 36. When the backup window period expires, the idle period begins (step 92). During the idle period, the non-compressed data stored on the hard drives 36 is retrieved (step 94) and provided to the micro-controller 50 through the bridge circuit 56, USB hub 54, and USB controller 52. The micro-controller 50 compresses the data (step 96) and then re-stores it on the hard drives 36 (step 98) to reclaim space on the hard drives 36. When the idle period is over and next backup window begins, a new duty cycle begins (step 84) and the aforementioned steps are repeated. In various embodiments, any one of a variety of software compression algorithms may be used, such as a zip; a gunzip; a bzip; a b2zip; a Lempel Ziv; and a LZS (Lempel Ziv Stac). Alternately, other compression algorithms can be used.

When compressed data on the emulated tape drive 20 is needed, it is retrieved and provided to the micro-controller 50. The data is decompressed using the software algorithms stored in the Flash memory 68 and provided to

the primary storage location 18 through appropriate input/output port 44. Since compression algorithms are typically asymmetric, data decompression is not nearly as computationally intensive as compression, the performance of the emulated tape drive 20 during data retrieval is not significantly degraded using a software solution.

The present invention thus provides an emulated tape device used for the backup of archival data that uses a software based data compression algorithm. Since the compression occurs when the emulated tape drive 20 is idle, the rate at which data is stored during a backup operation is not adversely effected in anyway.

It should be noted that the duty cycle could occur any time there is a detection of inactivity on the primary data interface 46. After some time period of inactivity (for example, 20 minutes), the system could begin retrieving and compressing data. This process can be interrupted at any point if activity is detected on the primary data interface 46. It is acceptable for data to be partially compressed, and it is possible to restart the compression from the point at which it was previously suspended.

Although the foregoing invention has been described in some detail for purposes of clarity of understanding, it will be apparent that certain changes and modifications may be practiced within the scope of the appended claims. For instance, the present invention can be ready practice with any type of read-write storage medium such as magnetic tape, silicon memory chips, devices such as SRAM, DRAM, Flash, EPROM, EUVPROM, EEPROM, etc. The described embodiments should therefore be taken as illustrative and not restrictive, and the invention should not be limited to the details given herein but should be defined by the following claims and their full scope of equivalents.